



How Sandia May Reach 1000 TFLOPS

Erik P. DeBenedictis

SAND 2003-2869P



Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company,
for the United States Department of Energy under contract DE-AC04-94AL85000.





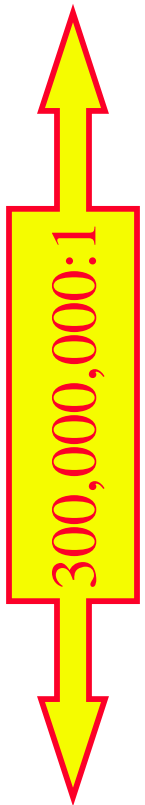
How To Reach 1000 TFLOPS?

- **Time Acceleration**

- **Scaling is mathematical based on # CPUs**
- **Time is physical based on clock rate**
- **If you keep architecture and # CPUs the same but increase the clock rate, the speed goes up, and efficiency stays the same**
 - **Scaling has to be right**

- **Has It Been Done Before?**

- **Cosmic Cube, 1981**
 - **64×50 KFLOPS**
- **nCUBE 10, 1988**
 - **1024×1 MFLOPS**
- **nCUBE 2, 1990**
- **Paragon, 1995**
- **ASCI Red, 1998-**
 - **9960×230 GFLOPS**
- **ASCI Red Storm, 2004-**
 - **10368×4 GFLOPS**
- **Petaflop?**





How to Spec the Machine?

- **If the Government sector specifies the machine**
 - It will be a linear speedup over ASCI Red
 - We will be able to predict performance
 - Project management will cost a bundle
- **If industry designs the machine**
 - It will have creative improvements designed to improve commercial potential
 - Untested improvements introduce risk
 - We are unlikely to be able to predict performance



Red Storm Scaling to 1000 TFLOPS

- Peak FLOPS 40 T \rightarrow 1000 T (25x)
- Per node 4 GFLOPS \rightarrow 100 GFLOPS
- Memory capacity stays about same at 1 byte/FLOPS
- Memory bandwidth 4 bytes/FLOPS stays the same
- Network bandwidth 4 bytes/FLOPS stays the same
- Latency (local/global) $2\ \mu\text{s}/5\ \mu\text{s} \rightarrow 80\ \text{ns}/200\ \text{ns}$

Ug

Much more Later

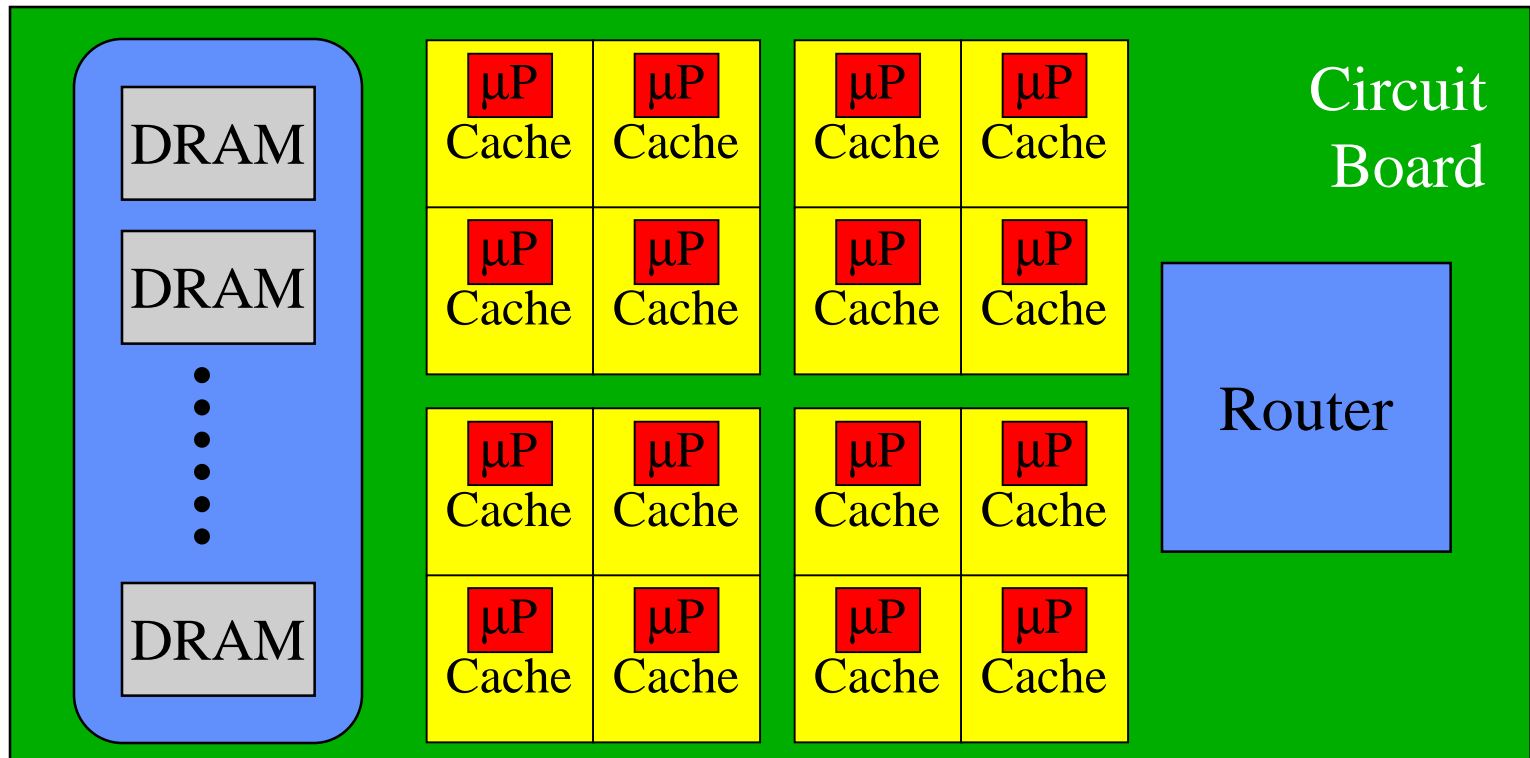
- Risk Factor
 - 100 GFLOPS CPU must be a SMP because cores cannot run this fast
 - However, various SMP nodes work OK up to n=8-16 (ASCI White, etc.)

Balance: time to operate on a number \sim time to send number across machine



Processors

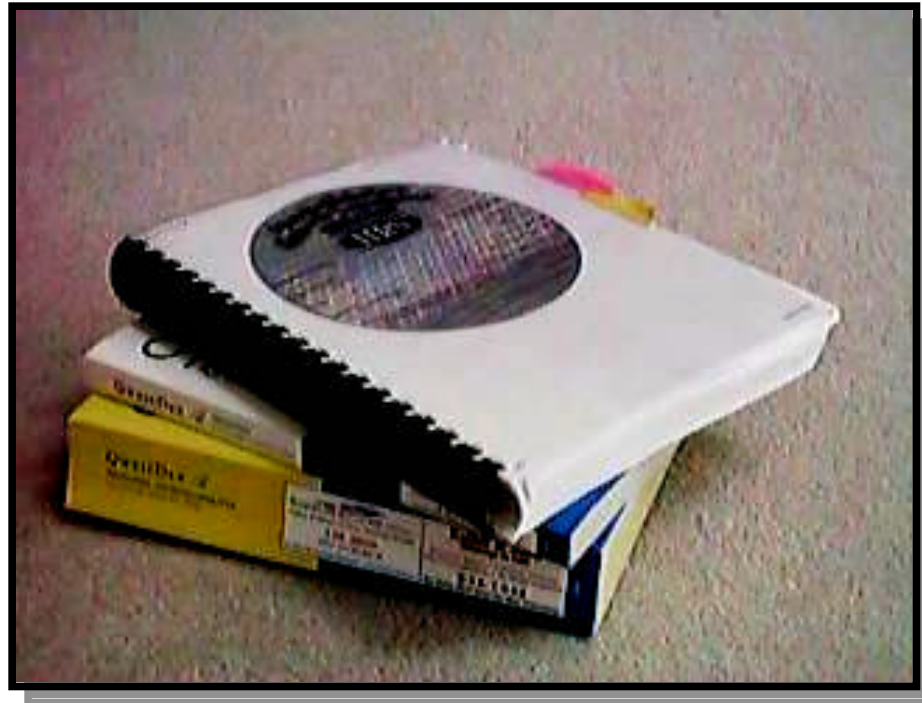
- **Fred Weber (AMD) @ Salishan**
 - “2008 144 GFLOP 4P * 4 * 9 GHz”



SIA Semiconductor Roadmap

- **Generalization of Moore's Law**

- Projects many parameters
- Years through 2016
- Includes justification
- Panel of experts
 - known to be wrong at times
- Size between Albuquerque white and yellow pages



Projected Interconnect Bandwidth

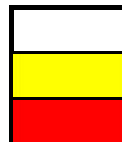
Table 23a High Frequency Serial Communications Test Requirements—Near-term

<i>Year of Production</i>		<i>2001</i>	<i>2002</i>	<i>2003</i>	<i>2004</i>	<i>2005</i>	<i>2006</i>	<i>2007</i>
<i>DRAM ½ Pitch (nm)</i>		<i>130</i>	<i>115</i>	<i>100</i>	<i>90</i>	<i>80</i>	<i>70</i>	<i>65</i>
<i>MPU / ASIC ½ Pitch (nm)</i>		<i>150</i>	<i>130</i>	<i>107</i>	<i>90</i>	<i>80</i>	<i>70</i>	<i>65</i>
<i>MPU Printed Gate Length (nm)</i>		<i>90</i>	<i>75</i>	<i>65</i>	<i>53</i>	<i>45</i>	<i>40</i>	<i>35</i>
<i>MPU Physical Gate Length (nm)</i>		<i>65</i>	<i>53</i>	<i>45</i>	<i>37</i>	<i>32</i>	<i>28</i>	<i>25</i>
<i>High-performance-level serial transceivers</i>								
Serial data rate (Gbits/s)		10	10	40	40	40	40	40
Maximum reference clock speed (MHz)		667	667	2500	2500	2500	2500	2500
<i>High-integration-level backplane and computer I/O</i>								
Serial data rate (Gbits/s) Production		2.5	3.125	3.125	10	10	40	40
Was	Introduction	3.125	—	10	—	40	—	—
Is	Introduction	3.125	—	10	—	40	—	—
Maximum port count at Production frequencies		20	100	200	100	200	100	200
at Introduction frequencies		—	—	20	—	20	—	—
Maximum reference clock speed (MHz) Production		166	166	166	667	667	2500	2500
Introduction		—	—	667	—	2500	—	—

White—Manufacturable Solutions Exist, and Are Being Optimized

Yellow—Manufacturable Solutions are Known

Red—Manufacturable Solutions are NOT Known



Pin Count

Table 3a Performance of Packaged Chips: Number of Pads and Pins—Near-term Years

<i>Year of Production</i>	<i>2001</i>	<i>2002</i>	<i>2003</i>	<i>2004</i>	<i>2005</i>	<i>2006</i>	<i>2007</i>
<i>DRAM ½ Pitch (nm)</i>	130	115	100	90	80	70	65
<i>MPU/ASIC ½ Pitch (nm)</i>	150	130	107	90	80	70	65
<i>MPU Printed Gate Length (nm)</i>	90	75	65	53	45	40	35
<i>MPU Physical Gate Length (nm)</i>	65	53	45	37	32	28	25
<i>Number of Chip I/Os (Number of Total Chip Pads) — Maximum</i>							
<i>Total pads—MPU</i>	3072	3072	3072	3072	3072	3072	3072
<i>Signal I/O—MPU (1/3 of total pads)</i>	1024	1024	1024	1024	1024	1024	1024
<i>Power and ground pads—MPU (2/3 of total pads)</i>	2048	2048	2048	2048	2048	2048	2048
<i>Total pads—ASIC high-performance</i>	3000	3200	3400	3600	4000	4200	4400
<i>Signal I/O pads—ASIC high-performance</i>	1500	1600	1700	1800	2000	2100	2200
<i>Power and ground pads—ASIC high-performance (½ of total pads)</i>	1500	1600	1700	1800	2000	2100	2200
<i>Number of Total Package Pins—Maximum [1]</i>							
<i>Microprocessor/controller, cost-performance</i>	480–1,200	480–1320	500–1452	500–1600	550–1760	550–1936	600–2140
<i>Microprocessor/controller, high-performance</i>	1200	1320	1452	1,600	1,760	1,936	2,140
<i>ASIC (high-performance)</i>	1700	1870	2057	2263	2489	2738	3012

Notes for Tables 3a and 3b:

[1] Pin counts will be limited for some applications where fine pitch array interconnect is used by PWB technology and system cost.

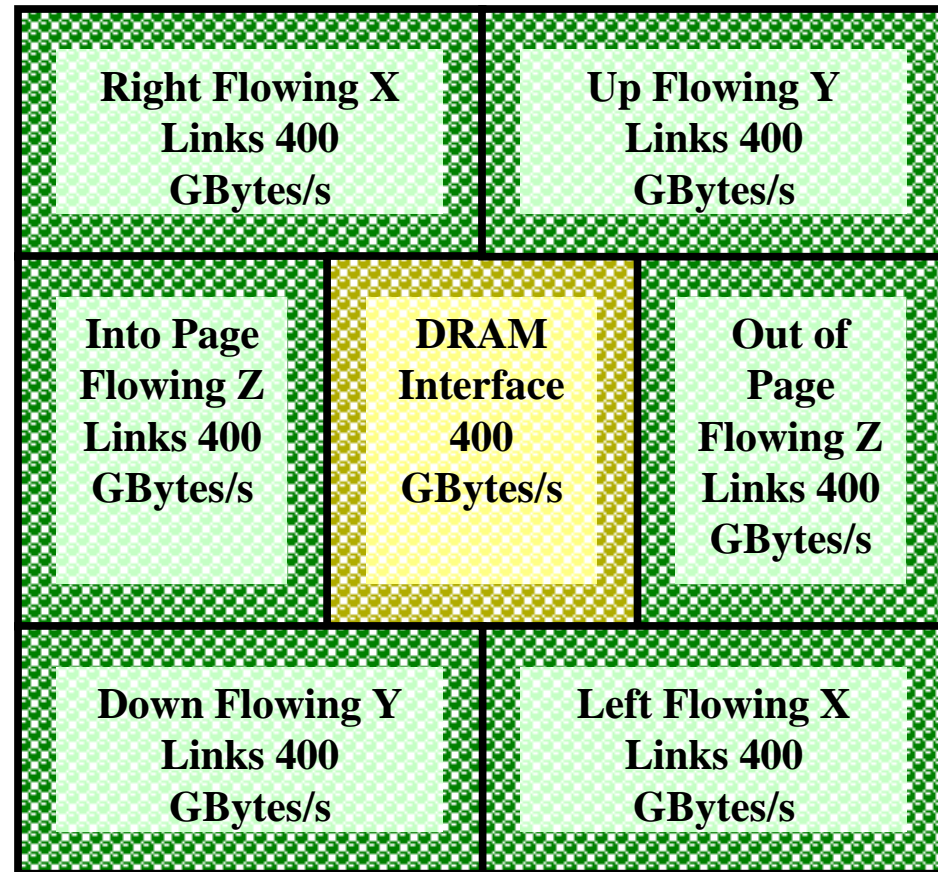
The highest pin count applications will as a result use larger pitches and larger package sizes.

The reference to signal pin ratio will also vary greatly dependent on applications with an expected range from 2:1 to 1:4



Chip Interconnect

- Bandwidth ought to be OK for next generation
 - Processor-Memory
 - Processor-Interconnect
- Remarkably unclear that COTS chips will exploit potential
 - This is a risk factor
- Diagram shows approximate proposed chip interconnect budget
- Bumps represent off-chip connections





Key Issue: Latency

- Of all semiconductor parameters, the speed of light (c) has fallen behind Moore's Law more than all others
 - c has not changed measurably in the last 30 years
 - c is decreasing exponentially with time when measured in distance traveled per clock period

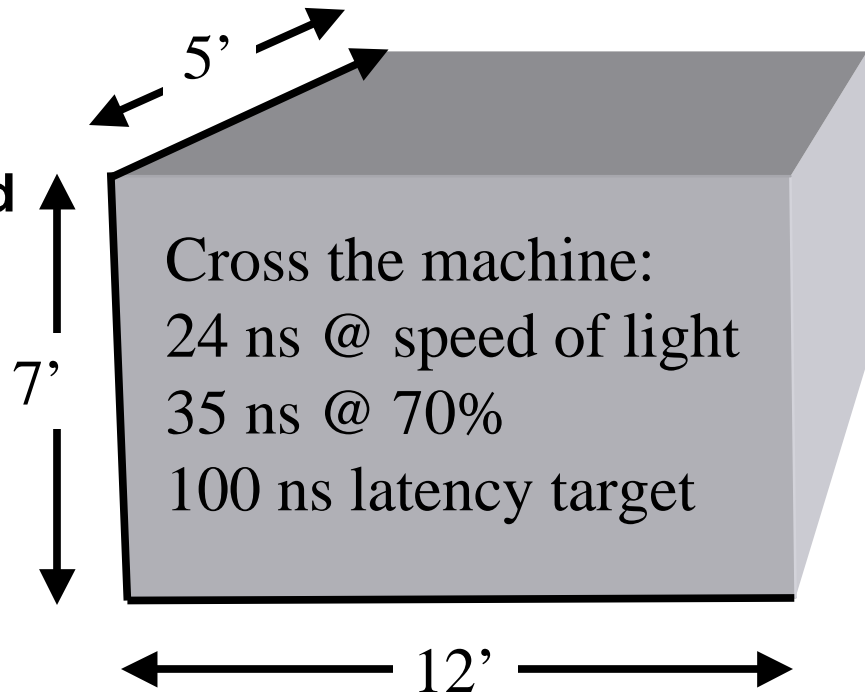


Joke

- Options
 - Machine is so large that 100 ns latency is not possible due to c .
 - Relax constraint
 - Not a good options because application scalability unknown with unbalanced latency
 - Cut size of machine
 - 3D packaging

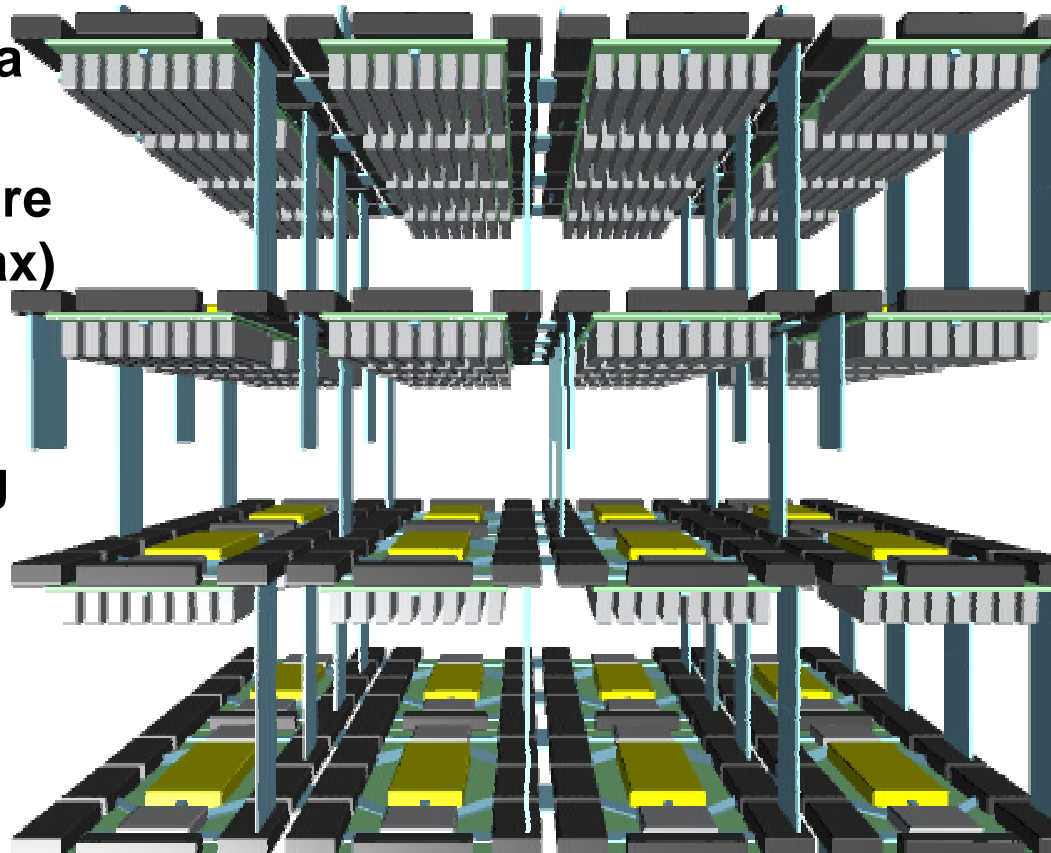
Cut Size of Machine

- **Water Cooling**
 - 10,368 nodes
(16×27×24)
 - Diagram to right would only be possible with water cooling
 - 100 ns latency
- **Air Cooling later**
- **About 10 ns budget for MPI software stack overhead**
 - (Need to talk to Barney)

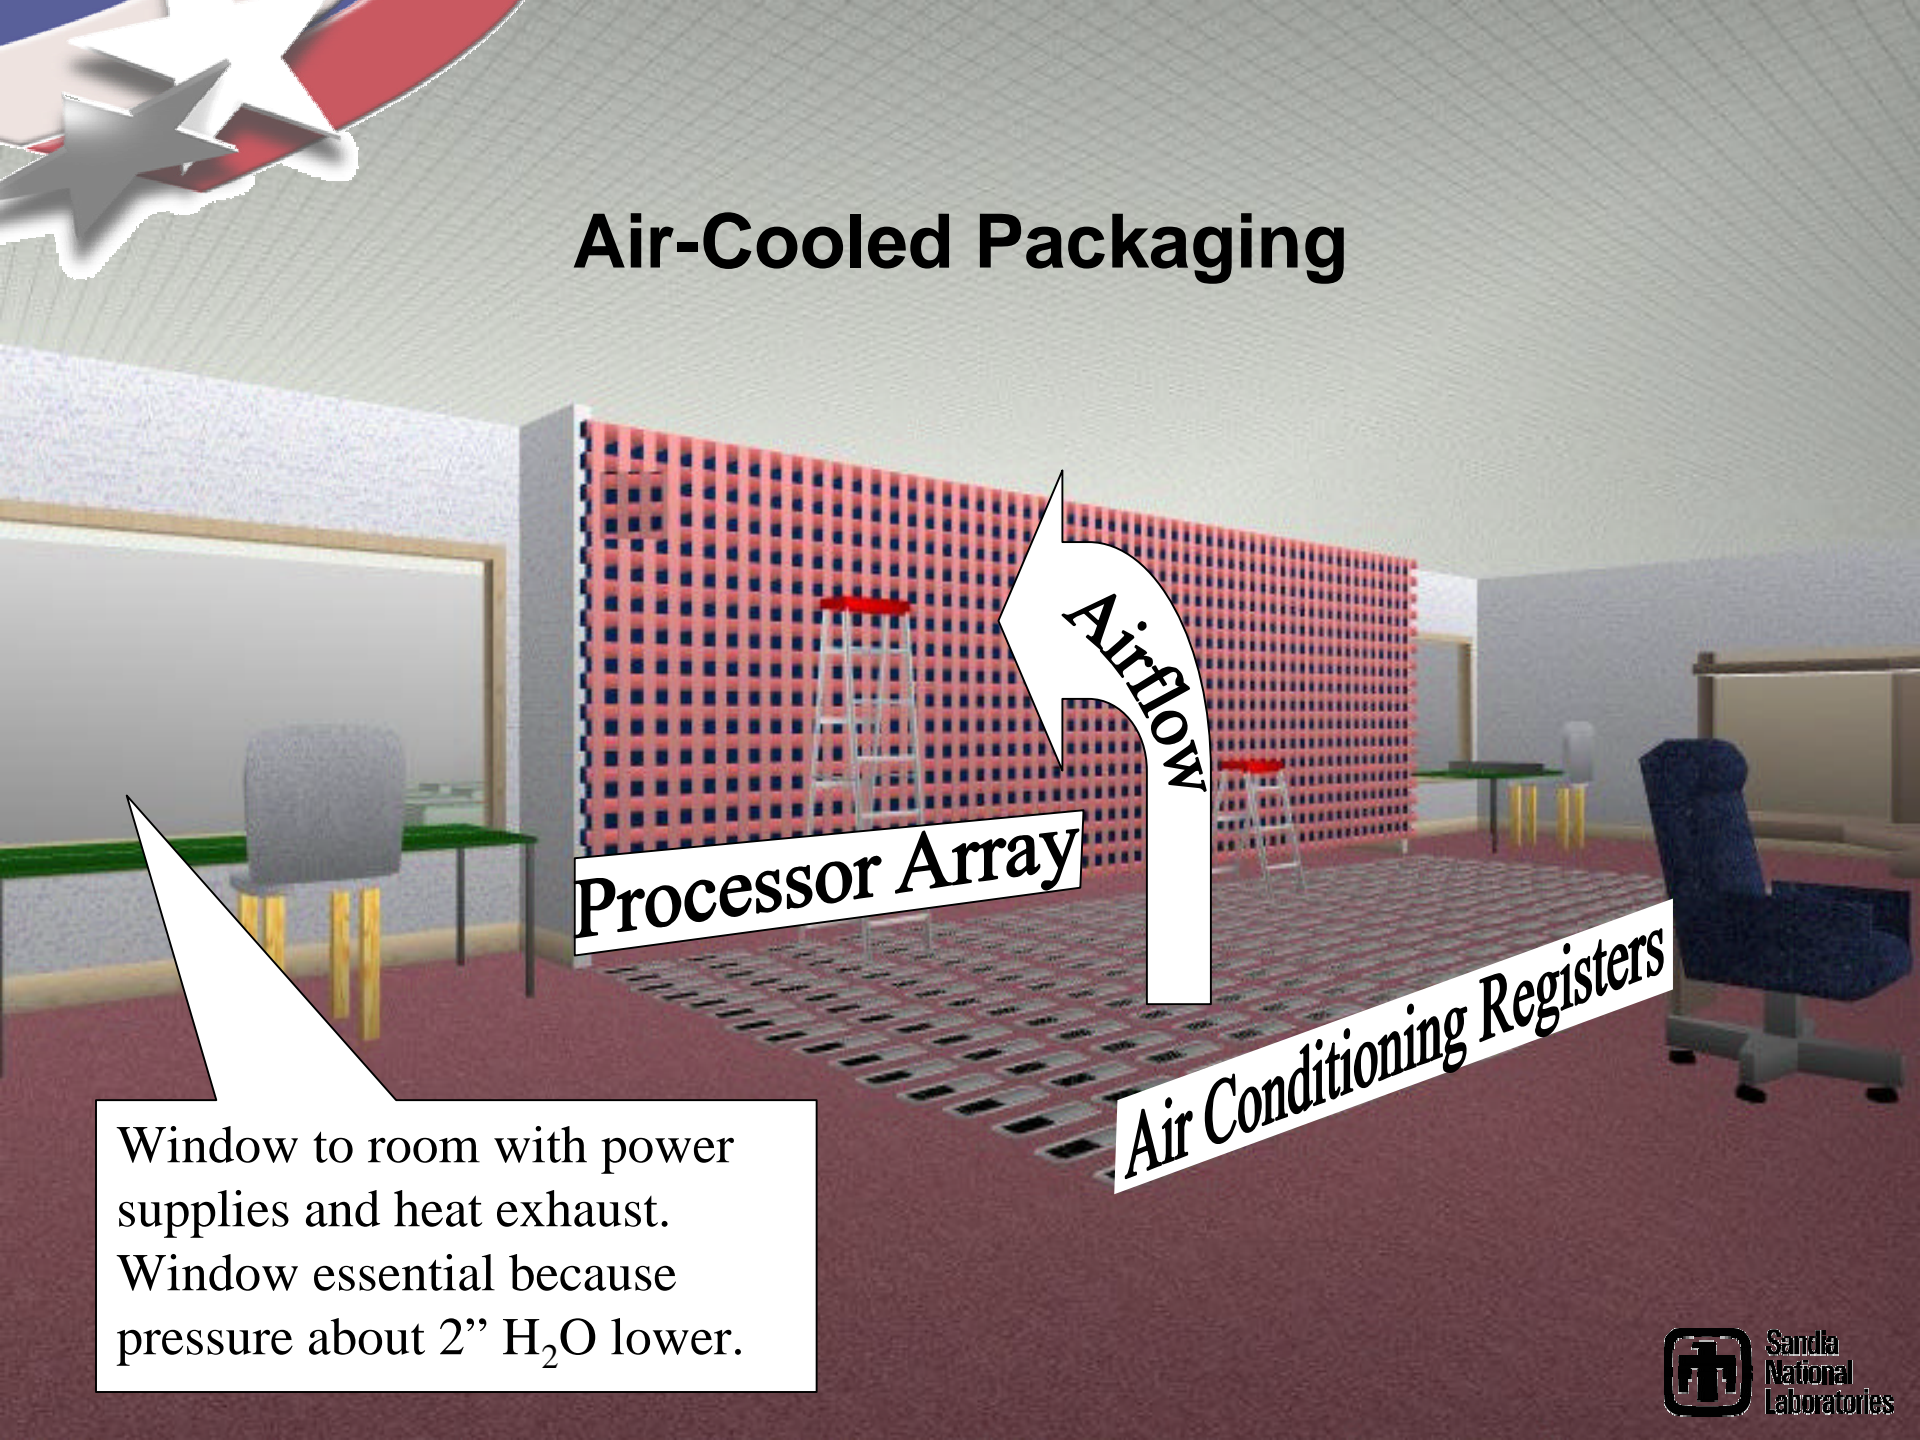


Homogeneous Packaging

- Entire supercomputer is a single structure
- All mesh network wires are of constant length (8" max)
- Air flows front to back
 - General approach will work for liquid cooling as well



Air-Cooled Packaging



Processor Array

Airflow

Air Conditioning Registers

Window to room with power supplies and heat exhaust. Window essential because pressure about 2" H₂O lower.



Is A Mesh A Good Topology?

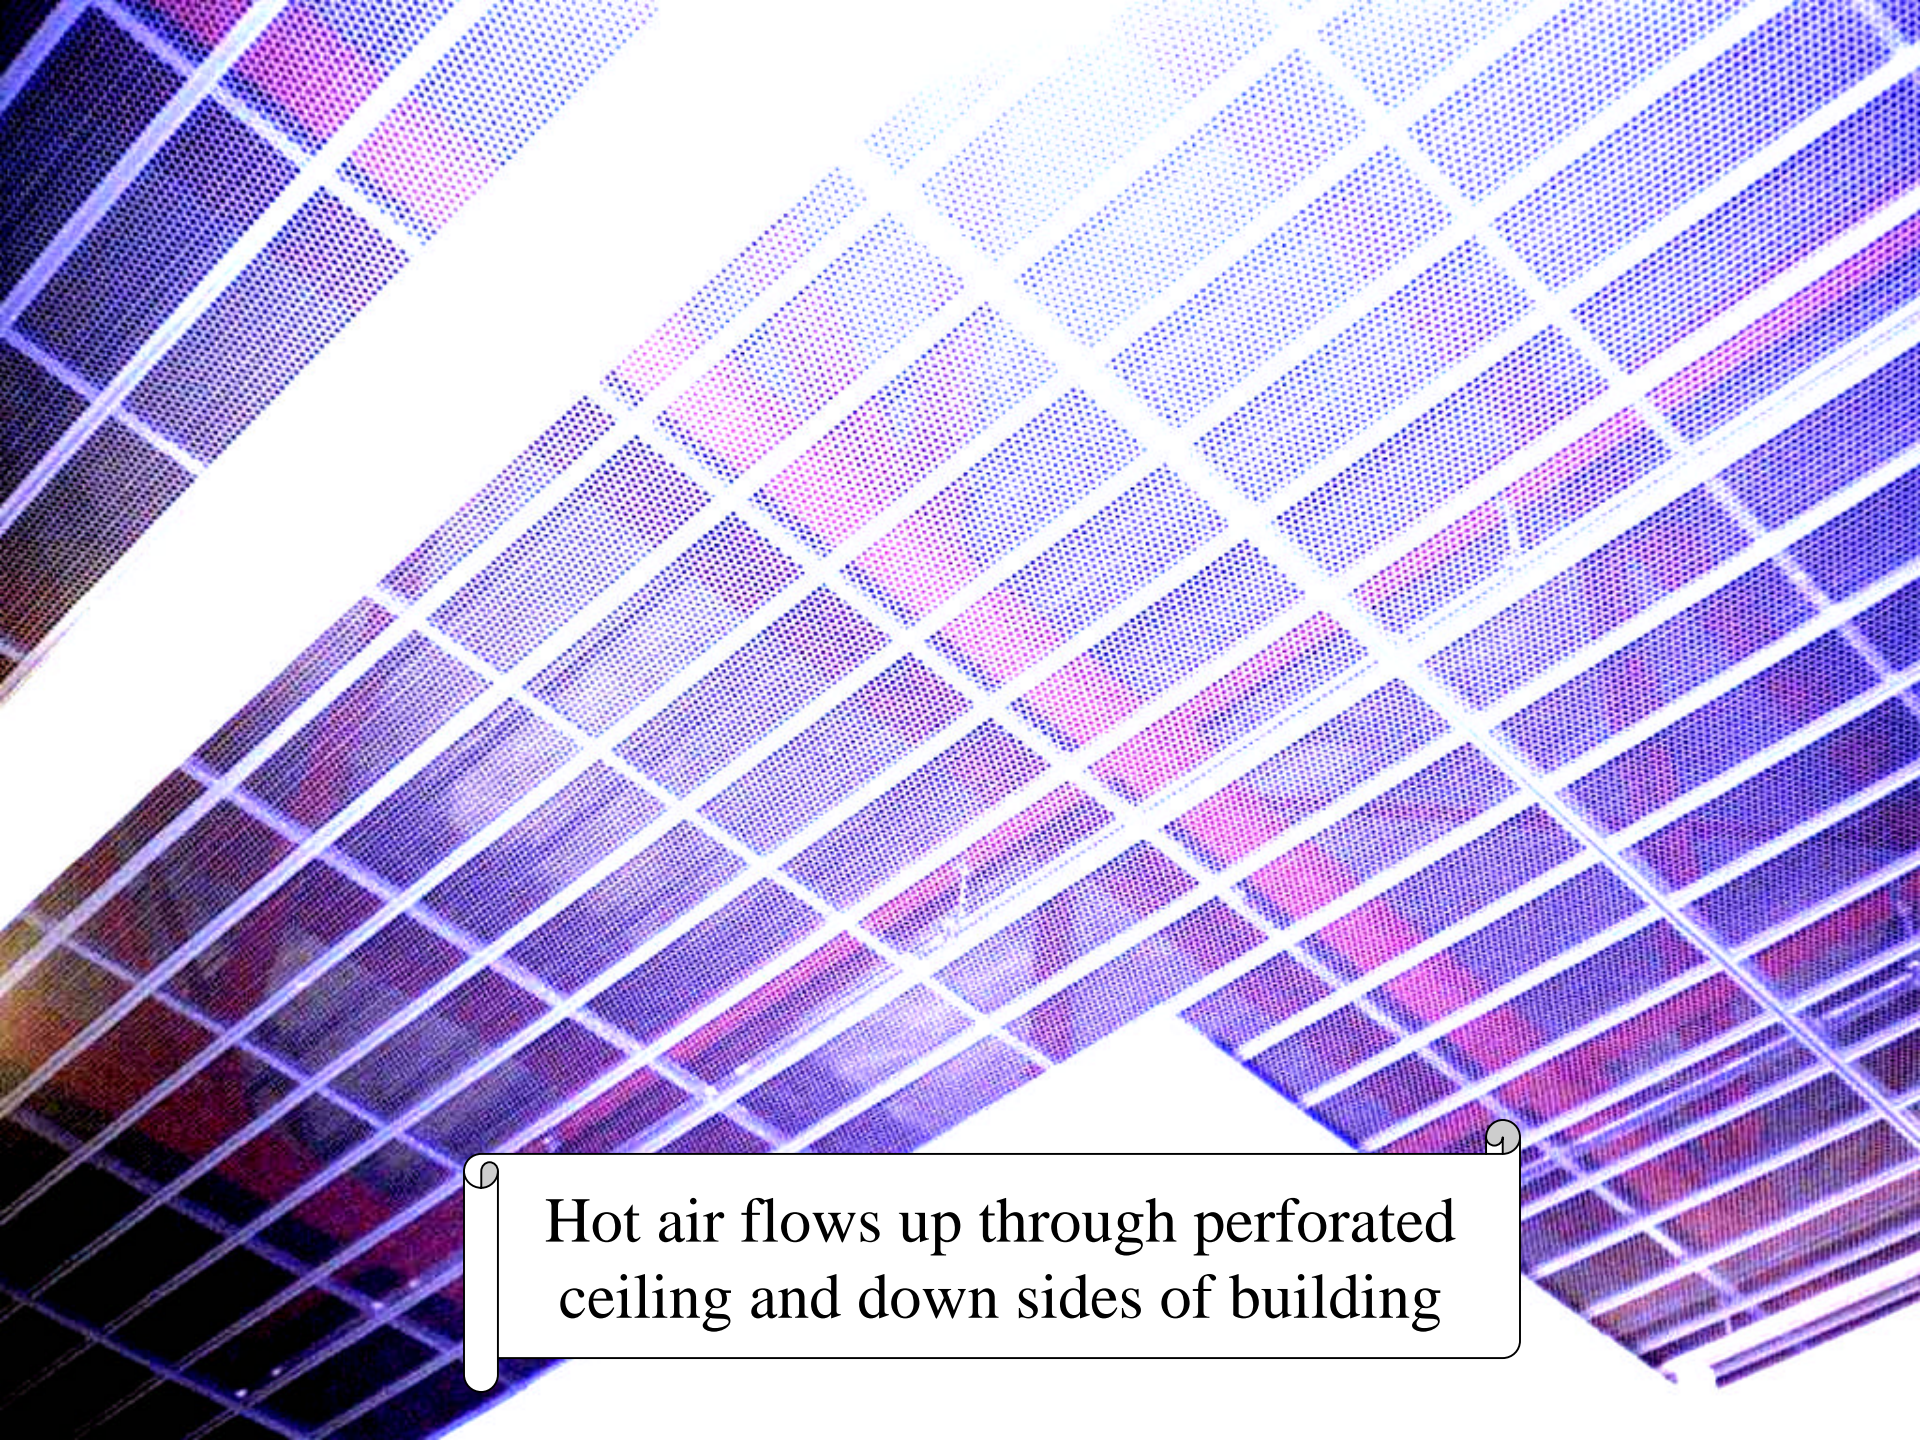
- **Mathematicians**
 - Delay related to number of “hops” or network diameter
 - Not relevant
- **Physicist**
 - Delay is distance traveled/c
- **Speed of propagation in proposed mesh is c divided by**
 - $\div \sqrt{3}$ Cartesian motion
 - \times propagation velocity in a transmission line (.7)
 - \div curvature of wire (2)
 - $+$ router delay (1 ns/hop)
- **Still within a constant factor of optimal!**



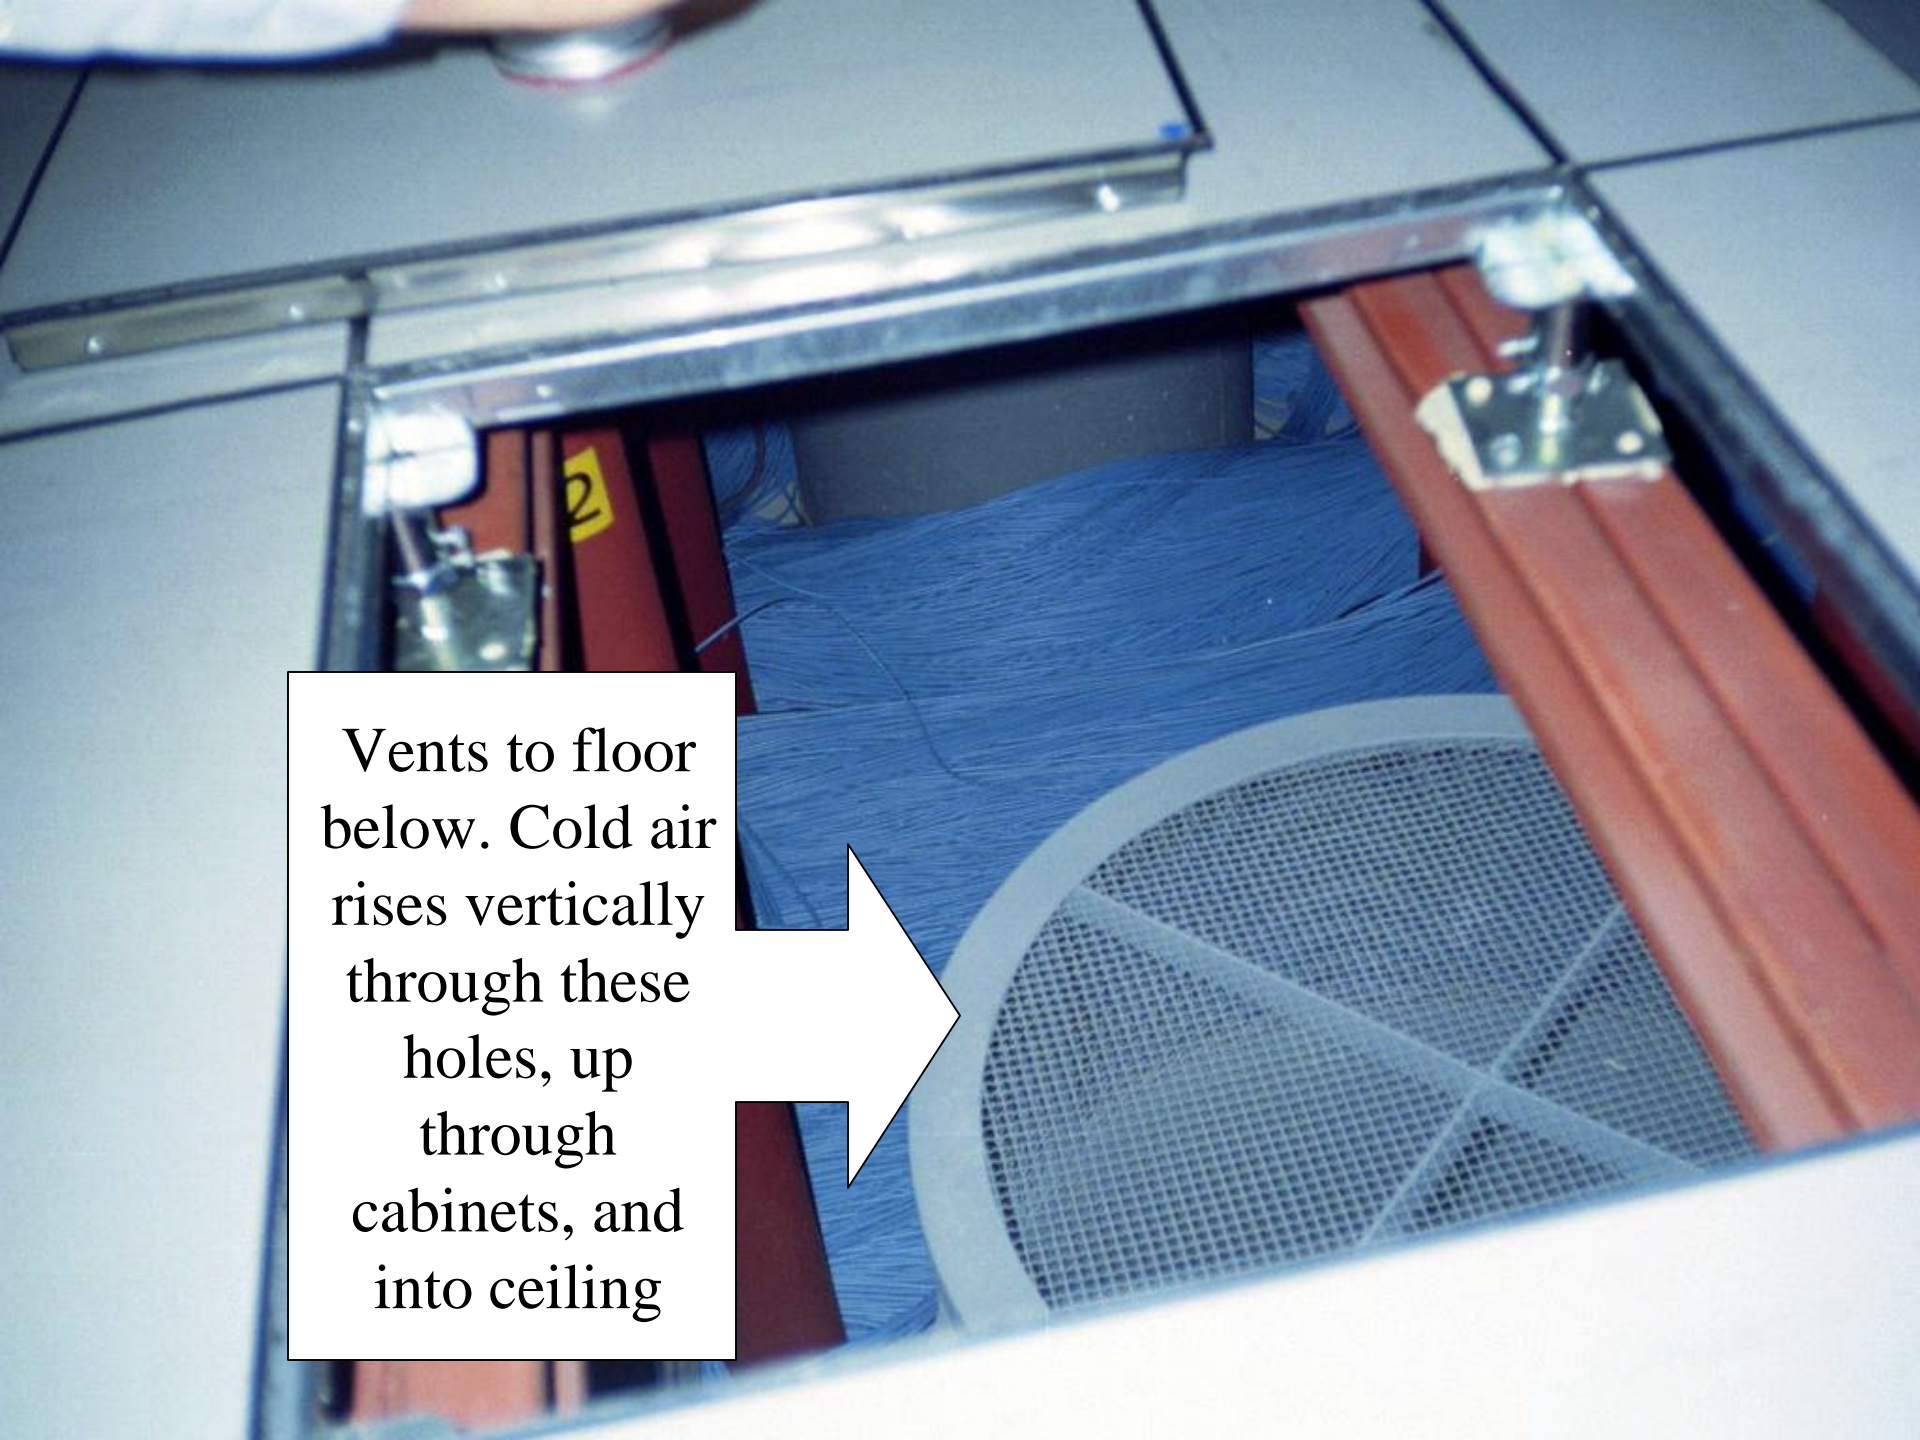
Has Anybody Made a 3d Machine?

- **All sorts of university prototypes**
- **Idea would be more credible if there were a successful example**
- **Let's see...**





Hot air flows up through perforated ceiling and down sides of building



Vents to floor
below. Cold air
rises vertically
through these
holes, up
through
cabinets, and
into ceiling



Reliability

- **Red Storm**
 - **Separate RAS network (2500 Unix processors & LAN)**
 - **Central point of information collection and control of entire machine**
 - **Capable of halting running machine, permitting deconfiguration of a faulty node, restart**
 - **Red Storm uptime specs: 50 & 100 MTBF/MTBI**
 - **If your PC had this MTBF, you'd take it back to Frys**



Reliability Forward

- **Cosmic Rays**

- These will be a problem in the next generation. COTS microprocessors have some tens of thousands of unprotected flip flops. This impacts an ASCI size machine on 6 month timeframe

- **FIT Rate**

- Manufacturers can give (under NDA) fit rate for components when used in a commercial environment
- Predicts ~1 hour MTBF
- However, machine rooms are temperature controlled and power is not cycled
- Actual “weeks” MTBF



Conclusions

- A 1000 TFLOPS successor to Red Storm is an engineering challenge
- Risk factors
 - SMP nodes
 - Memory bandwidth
 - Need new network interface
- Will 10 PFLOPS?
 - Scaling, speed of light, memory wall, threads
- COTS vs. Custom
 - Unknown which will “win”
 - Prepare to deal with both
 - Have capability on all key hardware issues
 - HDL
 - FPGA
- Shared address space?
 - 100 ns network makes this possible